LS-MolGen: Ligand-and-Structure Dual-driven Deep Reinforcement Learning for Target-specific Molecular Generation Improves Binding Affinity and Novelty

Song Li
Shanghai Jiao Tong University
Chao Hu
East China University of Science and Technology
Song Ke
Shanghai Matwings Technology Co., Ltd.
Chenxing Yang
Shanghai Matwings Technology Co., Ltd.
Jun Chen
Shanghai Matwings Technology Co., Ltd.
Yi Xiong
Shanghai Jiao Tong University
Hao Liu
Shanghai Jiao Tong University
Liang Hong (hongl3liang@sjtu.edu.cn)
Shanghai Jiao Tong University

Method Article

Keywords: molecule generation, transfer learning, reinforcement learning, ligand-and-structure-based drug design

Posted Date: April 12th, 2023

DOI: https://doi.org/10.21203/rs.3.rs-2793302/v1

Additional Declarations: Competing interest reported. Song Ke, Chenxing Yang, and Jun Chen are employees of Shanghai Matwings Technology Co., Ltd., Shanghai. Other authors declare no competing

interests.

METHODOLOGY

LS-MolGen: Ligand-and-Structure Dual-driven Deep Reinforcement Learning for Target-specific Molecular Generation Improves Binding Affinity and Novelty

Song Li[1], Chao Hu[4], Song Ke[6], Chenxing Yang[6], Jun Chen[6], Yi Xiong[3], Hao Liu[2,6] and Liang Hong[1,2,3,5,6,7*]

*Correspondence:
hongl3liang@sjtu.edu.cn
[1]School of Physics and Astronomy, Shanghai Jiao Tong University, Shanghai, China
[2]Institute of Natural Sciences, Shanghai Jiao Tong University, Shanghai, China
[3]School of Life Sciences and Biotechnology, Shanghai Jiao Tong University, Shanghai, China
[5]School of Pharmacy, Shanghai Jiao Tong University, Shanghai, China
[6]Shanghai Matwings Technology Co., Ltd., Shanghai, China
[7]Shanghai Artificial Intelligence Laboratory, Shanghai, China
Full list of author information is available at the end of the article

Abstract

Molecule generative models based on deep learning have attracted significant attention in de novo drug design. However, most current generative approaches are either only ligand-based or only structure-based, which do not leverage the complementary knowledge from ligands and the structure of binding target. In this work, we proposed a new ligand and structure combined molecular generative model, LS-MolGen, that integrates representation learning, transfer learning, and reinforcement learning. Focus knowledge from transfer learning and special explore strategy in reinforcement learning enables LS-MolGen to generate novel and active molecules efficiently. The results of evaluation using EGFR and case study of inhibitor design for SARS-CoV-2 Mpro showed that LS-MolGen outperformed other state-of-the-art ligand-based or structure-based generative models and was capable of de novo designing promising compounds with novel scaffold and high binding affinity. Thus, we recommend that this proof-of-concept ligand-and-structure-based generative model will provide a promising new tool for target-specific molecular generation and drug design.

Keywords: molecule generation; transfer learning; reinforcement learning; ligand-and-structure-based drug design

Introduction

The main goal of de novo drug design for a specific target is to identify novel active molecules that can simultaneously satisfy certain desirable properties, such as the generated molecules should have similar or better binding affinity and be chemically diverse to already known ligands [1]. In recent years, various artificial intelligence (AI)-based molecular generative models have been proposed to address such needs since their potential for chemical space exploration [2, 3, 4, 5, 6, 7, 8, 9]. Most of these generative approaches are either ligand-based [5, 10, 11, 12, 13, 14, 15] or structure-based [16, 17, 18, 19, 20, 21]. For example, for ligand-based methods, VAE [11] used a variational autoencoder to convert discrete representations of molecules to a multidimensional continuous representation, and generates new molecules for efficient exploration and optimization through open-ended spaces of known ligands. latentGAN [12] generates target-biased compounds through combining an autoencoder and a generative adversarial neural network. AAE [10] combined the idea of

VAE with that of adversarial training as found in GAN and was applied to the molecule generation task. In contrast, for structure-based methods, REINVENT [2] provided a running mode for goal-directed molecular generation by integrating a reinforcement learning. Pocket2Mol [21] proposed a 3D generative model with an auto-regressive sampling scheme that generates molecules given a designated 3D protein binding site. MolDQN [19] leverages domain knowledge of chemistry and reinforcement learning to optimize molecule without pre-training on any dataset and only with a target protein structure. Several methods also combined the knowledge from both ligands and the target structure, such as SBMolGen [22] integrates an RNN generative model with a Monte Carlo tree search (MCTS), and docking simulations. OptiMol [18] combines the graph to Selfies VAE, a docking program, and the clamped version of Conditioning by Adaptive Sampling (CbAS) to achieve binding affinity optimization.

In general, the ligand-based approach requires sufficient known ligands, whose bioactivity has been proved in cellular, animal, or even clinical experiments as reference to generate new molecules with optimal properties as shown in Figure 1a. The structure-based method tends to generate molecules by capturing and optimizing the physical interaction of ligands with the structure of the target protein as Figure 1b depicted. However, the ligand-based method could be biased towards the reference ligands in the training set and the generated molecules resemble them with low structural diversity (collapse to the yellow region in illustration of (Figure 1a), and the structure-based method may be hard to efficiently converge to chemical space with sufficiently high bioactivity, particularly when examined by cellular or animal experiments [23]. As a more inclusive strategy, ligand-and-structure-based method (Figure 1c) leverages the complementary information of known ligands and target protein structure, generating novel compounds that chemically cover the known bioactive compounds, and even with improved binding affinity.

In this study, we constructed a benchmarking dataset from the latest version of ChEMBL and proposed a ligand-and-structure dual-driven deep reinforcement learning method (LS-MolGen) for target-specific molecular generation. By making use of the complementary knowledge of known bioactive ligands and the structure of the specific target, and integrating docking score and a diversity filter in the inception of the reinforcement learning, LS-MolGen exhibited comparable performance to various state-of-the-art molecule generative models as demonstrated by the testing on the EGFR and Mpro. The results show that the generated molecules have better novelty, diversity, and affinity. A further analysis of molecular structure and molecular docking of Mpro inhibitors has confirmed that LS-MolGen has the capability to generate both novel and active compounds.

Methods and materials

Data preparation

Above 1.9 million small molecules represented in canonical SMILES format were downloaded from the ChEMBL database (version 31) [24]. For all SMILES strings, data preprocessing is necessary, including standardization of charge, removing small fragments and metals, removing duplicates and invalid SMILES. Filtered by molecular weight in the range from 150 to 650 Daltons, molecular heavy atoms within

C, N, O, S, F, Cl, Br, and filtered via medicinal chemistry filters (MCFs) and PAINS filters, a benchmarking dataset with 1.5 million molecular structures were constructed.

In this work, we selected human epidermal growth factor receptor (EGFR, PDB ID: 2RGP, Uniprot ID: P00533) and the SARS-CoV-2 main protease (Mpro, PDB ID: 7L12, Uniprot ID: P0DTD1) as two test targets. 2005 ligands which had binding bioactivity within 100 nM, i.e., pChEMBL value (pX, including pKi, pKd, pIC_{50}, or pEC_{50}) no smaller than 7 toward the EGFR target were extracted from the benchmarking dataset. And 127 inhibitors targeting Mpro were extracted from the reported paper [25] and the benchmarking dataset with pIC_{50} no smaller than 5. These known ligands will be used for fine-tuning the generative models in the transfer learning process.

Architecture of LS-MolGen method

A brief overview of the entire architecture of LS-MolGen model is illustrated in Figure 1d, which consists of three essential sub-models. The first model is an initial-stage model for pre-training of the general knowledge called the prior model (Figure 1e), based on the recurrent neural network (RNN). The second model is a transfer model (Figure 1f), based on the same architecture as the prior model, for transfer learning of the general knowledge to the focused knowledge via sharing the prior networks and reweighting the layers. The last one is an agent model (Figure 1g) fine-tuned by reinforcement learning guided by molecular docking for the structure-based rational exploration of chemical space. Here, the docking score that estimates the strength of the interaction between target and ligand is obtained from LeDock (`http://www.lephar.com`), which exhibits an outstanding performance in a recent comprehensive evaluation of docking programs on a diverse set of 2002 protein-ligand complexes [26]. Hence, LS-MolGen is a ligand-and-structure dual-driven model which makes full use of the knowledge of both known active ligands and protein structure to generate molecules with high binding affinity and novelty.

Pre-training

Each molecule in benchmarking dataset represented in SMILES format was split into a series of tokens and then all tokens were collected to construct the SMILES vocabulary, resulting in final vocabulary containing 101 tokens. Two additional tokens, GO and EOS, may be added to denote the beginning and end of a SMILES respectively. Novel valid SMILES sequences could be generated through property grammar which was learned from numerate SMILES in benchmarking dataset.

The prior model is consist of a standard RNN model, whose architecture is provided in Fig. 1e. It contains six layers: one input layer, one embedding layer, three recurrent layers, and one output layer. Molecules, represented by a sequence of tokens, can be received as features by the input layer. In the embedding layer, each token was encoded into a 128-dimensional embedding vector. For a recurrent layer, a gated recurrent unit (GRU) [27] was used as the recurrent cell with 512 hidden neurons.

The output at each position was the probability that determined which token in the vocabulary would be chosen to grow the SMILES string. Maximum likelihood

estimation was employed as following loss function to train the RNN:

$$NLL\left(s\right) = -\sum_{i=1}^{n} \log P\left(x_i \mid x_{<i}\right) \tag{1}$$

Transfer learning

To guide the generative model toward areas with similar or high binding affinity to known ligands in the chemical space, we subjected it to a transfer learning scenario while aiming to learn the distribution of the bioactive ligands. It is achieved by directly transferring the parameters of the prior model [13], which is computationally cost-effective. This model requires a pre-trained prior model with a generative capacity and the potential to sample compounds from a rather vast chemical space. And the prior is subjected to transfer learning with a smaller set of known bioactive ligands. The loss function during transfer learning is the same as pre-training.

Reinforcement learning

We introduced molecular docking score as a component of scoring function into reinforcement learning to guide the agent model exploring novel sub-areas in the chemical space, for example, enhancing the ability of the agent for scaffold hopping that cannot be fully captured by the focused transfer model. The agent in RL is an exploration model which shares identical architecture and vocabulary with the prior model. Essentially, the agent was initialized with the pre-trained prior network at the beginning of the RL. We adopted the RL algorithm in REINVENT2.0 [7] to build a customized scoring function for the molecular optimization:

$$Score\left(A\right) = \begin{cases} \frac{\max\left(ds, k\right)}{k} & unsatisfy \quad DF \\ 0 & satisfy \quad DF \end{cases} \tag{2}$$

ds is the docking score. A scaffold similarity diversity filter (DF) [7] is used to evaluate whether the SMILES string has been sampled before or whether it satisfies the DF policy. The Score will be set to 0 if the DF filters determine that the sampled compound already exists or if there are too many compounds of the same scaffold and their number exceeds our predefined threshold. And k is a rescale parameter for scoring function normalized to [0, 1].

The scoring function is combined with the likelihood of transfer and used to form the augmented likelihood:

$$\log P\left(A\right)_{Augmented} = \log P\left(A\right)_{Transfer} + \sigma \cdot Score\left(A\right) \tag{3}$$

The scoring function is multiplied by σ which is a scalar coefficient used for scaling up the scoring function output to the same order of magnitude as the likelihood. Analogously, the loss is calculated as the squared difference between the likelihood of agent and the augmented likelihood:

$$loss = \left[\log P\left(A\right)_{Augmented} - \log P\left(A\right)_{Agent}\right]^2 \tag{4}$$

Thus, after the agent samples a batch of SMILES, the reward is contributed from several components: docking score, diversity filter, and the likelihood from transfer model.

Docking protocol and Pharmacophore matching

The co-crystal structures of the target-ligand complexes were downloaded from the RCSB Protein Data Bank (`https://www.rcsb.org/`). Protein preparation was processed by Lepro (`http://www.lephar.com/`), and LeDock program was utilized to conduct docking where five binding poses were generated and RMSD threshold was set at 1 Å. The web server of Pharmit [28] was used to match pharmacophores and select molecules. Eventually, PyMOL [29] was applied to draw figures.

Evaluation metrics for generative model

MOSES [30] is a main benchmark in the field of de novo molecular generation, but it may not quietly suitable for target-focused molecular generative model. Consequently, this study provides more acceptable benchmark metrics to access the overall quality of generated molecule dataset G based on training set T or known ligand set L, of course, some of which are borrowed from MOSES.

`Validity`: a measure of the proportion of generated molecular SMILES that can be successfully parsed and validated by RDKit, which is often used to evaluate the quality and correctness of generated molecules, and can help identify generation methods that produce chemically plausible structures.

`Uniqueness`: the proportion of the unique and valid structures generated.

`Diversity`: the proportion of unique scaffold of generated molecules which provides insight into the richness and variety of the generated structures.

`Novelty`: the proportion of unique scaffold of generated molecules that are not in test set T or the known ligand set L, which provides insight into the originality and distinctiveness of the generated molecules.

`Recovery`: the proportion of molecular fragments in the known ligand set L covered by fragments from generated set G, which should not be too high or too low. High recovery suggests lack of novelty, while low recovery indicates poor use of existing fragments.

`Active rate`: the proportion of generated molecules with high binding affinity (docking score $<$ -7 kcal/mol) against target protein, and can give the insight into the potential effectiveness of the generated structures as potential drug candidates. Here, the threshold of -7 kcal/mol refers to the cutoff pChEMBL value of known bioactive ligands introduced in the data preparation section.

`Success rate`: the proportion of generated molecules that both satisfy maximum similarity to molecules in T $<$ 0.7 and docking score $<$ -7 kcal/mol, which is considered to be a more practical metric to evaluate the potential efficacy and novelty of generated molecules.

1 Results and discussion

1.1 Model performance

We assembled these seven metrics: validity, uniqueness, diversity, novelty recovery, active rate, and success rate, to comprehensively assess model performance and

identify potential strengths and limitations compared with other previously proposed generative models [10, 11, 12, 2, 21, 19, 22], emphasizing the assessment of both the binding affinity and novelty. Taking the human epidermal growth factor receptor (EGFR) as the first test case, we re-trained all these baseline models and trained LS-MolGen to de novo design inhibitors of EGFR.

Referencing the proposed pipeline in MOSES [30], for the ligand-based generative model, we used the models and hyperparameters available in the platform, such as the AAE [10], VAE [11], and LatentGAN [12]. For structure-based model, REINVENT [2], Pocket2Mol [21], and MolDQN [19] were re-trained with the same hyperparameters reported in their papers. And for ligand-and-structure-based model, except the method proposed here, we re-trained another model, SBMolGen [22], for comparison. Training data was the benchmarking dataset constructed in this work and the collected known ligands.

For a fair comparison of their performance, we trained and sampled 5,000 molecules for six times independently from each model, and applied the evaluation metrics mentioned above for a comprehensive benchmark as shown in Table 1. The proposed LS-MolGen model performs preferentially over other generative models with similar or high percentage of validity and uniqueness SMILES strings. In terms of diversity, novelty, active rate, and success rate, the values of our method are better than all other methods. And compare to structure-based model such as MolDQN, the running time of generating 5000 molecules is much shorter (4.6 hours vs 2 days, see Table S1 in Supporting Information (SI)). The focus knowledge from transfer learning and special explore strategy (docking score and diversity filter) in reinforcement learning helped LS-MolGen exhibiting superior performance in efficiently generating new molecules with high affinity, large diversity, and novelty. As expected, ligand-based methods generate molecules with similar high affinity (high active rate) as the known bioactive ligands, but low novelty that most of their structures were similar to the known ligands, and structure-based methods generate molecules more diverse but still lacking in overall success rate.

Figure 2a indicates that the exploration by the Agent is working during the reinforcement learning process. The score is contributed from docking score, diversity filter, and the likelihood of known ligands, and the continuously rising score demonstrates the model's ability to effectively explore chemical space for novel and active molecules. To better understand the chemical space of the generated molecules, we evaluated the chemical space coverage by calculating the Morgan fingerprint [31] used as a t-distributed Stochastic Neighbor Embedding (t-SNE) visualization. As Figure 2b shown in the t-SNE plot, the generated molecules not only cover the chemical space with the known EGFR inhibitors but also extend it to the new chemical space. Figure 2c exhibits that the generated molecules have higher affinity compared to the known bioactive ligands. Furthermore, we compared the distribution of properties (molecular weight, QED, and SA score). As shown in Figure 2d-f, the properties distributions of generated molecules are close to that of known ligands. Considering the above quantitative evaluations, the model we developed, LS-MolGen, demonstrates satisfactory results by generating both novel and active molecules and outperforming other methods in the potential for drug design.

Ablation experiment for LS-MolGen

We conducted an ablation experiment to assess the performance of our proposed LS-MolGen framework in generating molecules targeting EGFR. The results of this experiment are summarized in Table 2, where the LS-MolGen framework is dissected into its key components—Prior, Prior+Transfer, Prior+Agent, and the full model (Prior+Transfer+Agent)—in order to evaluate the individual and combined contributions of these components on the overall performance.

In the ablation study, the comprehensive LS-MolGen model outperformed its individual components, achieving optimal scores for Validity and Active rate, and high scores for Uniqueness, Diversity, Novelty, and Success rate. These findings highlight the necessity of integrating all components within the LS-MolGen framework to enhance molecular generation, particularly the Agent model, which plays a pivotal role in promoting novelty and diversity. Accordingly, our investigation demonstrates the potential of the LS-MolGen model in furthering the progress of molecular design research.

Molecule design using LS-MolGen

`Generating potential inhibitors of SARS-CoV-2 main proteinase (Mpro).` To evaluate whether LS-MolGen generated compounds could be novel and potent molecular candidates, we used Mpro as the second test case. Mpro is a potential target for discovery of therapeutic agents for treatment of COVID-19. Several Mpro inhibitors were collected from ChEMBL [24] and the reported paper [25], and the co-crystal structure of Mpro complexed with an inhibitor is obtained from RCSB PDB (ID: 7L12). The goal was to generate more novel and potential inhibitors of Mpro. With these known bioactive inhibitors and the structure of Mpro, LS-MolGen was applied to generate 5,000 compounds, whose properties distributions and chemical space were illustrated in Figure S2. Compound 1 (PF-0835231) (IC_{50} = 8 nM) [32], compound 2 (IC_{50} = 30 nM) [24] and compound 3 (IC_{50} = 128 nM) [25] were revealed as potent inhibitors of SARS-CoV-2 Mpro. By matching the pharmacophore of the inhibitor, the most matching compound is selected from the generated molecules (compounds 4-6 in Figure 3). As it depicted, the similarity between inhibitors and the corresponding generated compounds is low, and the molecular scaffold is completely new, which realizes scaffold hopping. This result demonstrates that LS-MolGen is capable of generating potential molecules that differ from known bioactive inhibitors.

Figure 4 depicts the docked poses of compounds 4, 5, and 6, with docking scores of -9.71, -9.02, and -9.45 kcal/mol, respectively. In comparison, the redocking scores of compounds 1, 2, and 3 were -8.24, -7.12, and -9.34 kcal/mol, respectively. It was noted that the redocking scores of compounds 1 and 2 were much worse compared to their bioactivities due to their covalent inhibitor nature, which cannot be captured by molecular docking. Importantly, the generated molecules displayed more potent binding affinity compared to the known inhibitors, validated by both molecular docking and MM-GBSA (see details in SI), further demonstrating LS-MolGen's ability to generate molecular candidates with similar or improved binding affinity. In addition, the pockets occupied by the three molecules and their binding modes are largely consistent, which is in accordance with the known inhibitor binding patterns. It is noteworthy that while these three molecules have conserved interactions

with the Glu-166 residue in the pocket, the other interactions differ. For example, compound 4 interacts with Cys-145 in a manner similar to the covalent binding mechanism of PF-0835231, the interaction of compound 5 extends to the farther Gln-192 residue, and compound 6 forms a hydrogen bond with Phe-140. These findings highlight the ability of LS-MolGen to discover novel interactions with the target protein and leverage it to generate novel and potent compounds. Of course, a further experimental study is necessary to determine the potency of these designed compounds in inhibiting Mpro.

Conclusion

In this study, we have developed LS-MolGen, a novel molecular generative model for de novo drug design that generates molecules bases on the knowledge from both bioactive ligands and structure of target protein. LS-MolGen consists of three essential sub-modules, which are representation learning, transfer learning and reinforcement learning. Representation learning was responded for learning the general grammar rules of SMILES string, transfer learning focused on the chemical space of bioactive ligands, and reinforcement learning was used to explore the potential chemical space. The inception with docking score and a diversity filter in RL guided Agent to discover active and novel compounds.

In the case study of EGFR, our model shows better performance to several other state-of-the-art molecule generative model. The high values of metrics of novelty, diversity, active rate and success rate, as well as the chemical space analysis and the distribution of docking scores, demonstrates that LS-MolGen is of capable generating both novel and high-affinity molecules. We explored the potential inhibitors design of SARS-CoV-2 Mpro as another case study. The results show that LS-MolGen was able to de novo design promising compounds with novel scaffold and high molecular docking score, which captures important pharmacophore features of known inhibitors.

Considering the docking score that reflects the interactions between the ligand and the target was used as the reward score in the reinforcement learning, of course it can be changed to other scores, such as ADMET scores or similarity scores, to address a variety of design needs. To summarize, LS-MolGen proof the concept of ligand-and-structure-based generative model performance better and provide a promising new tool for target-specific molecular generation.

Acknowledgements

The authors thank the Center for High Performance Computing at Shanghai Jiao Tong University for computing resources.

Funding

This work was supported by the grants from the National Science Foundation of China [grant numbers 11504231, 21873101, 31630002, 32030063], the Innovation Program of Shanghai Municipal Education Commission [2019-01-07-00-02-E00076], and the student innovation center at SJTU.

Abbreviations

Abbreviations used in this article: AI, artificial intelligence; MCTS, Monte Carlo tree search; CbAS, Conditioning by Adaptive Sampling; MCFs, medicinal chemistry filters; EGFR, epidermal growth factor receptor; Mpro, the SARS-CoV-2 main protease; RNN, the recurrent neural network; GRU, gated recurrent unit; DF, diversity filter; SI, Supporting Information; t-SNE, t-distributed Stochastic Neighbor Embedding.

Availability of data and materials

The datasets and source code are publicly available at https://github.com/songleee/LS-MolGen. The Supporting Information is available free of charge at:
The hyper-parameters of Models and training, Supplementary Tables and Figures.

Competing interests

Song Ke, Chenxing Yang, and Jun Chen are employees of Shanghai Matwings Technology Co., Ltd., Shanghai. Other authors declare no competing interests.

Authors' contributions

L. H. and H. L. designed and supervised the project. S. L. and C. H. contributed equally in this work. S. L. designed and implemented the deep learning model and performed the model training. C. H. tested the baseline models. S. K., C. Y. and J. C. discussed and analyzed the data. S. L. wrote the manuscript. L.H., H. L. and Y. X. revised the manuscript.

Author details

[1]School of Physics and Astronomy, Shanghai Jiao Tong University, Shanghai, China. [2]Institute of Natural Sciences, Shanghai Jiao Tong University, Shanghai, China. [3]School of Life Sciences and Biotechnology, Shanghai Jiao Tong University, Shanghai, China. [4]Department of Computer Science and Engineerings, East China University of Science and Technology, Shanghai, China. [5]School of Pharmacy, Shanghai Jiao Tong University, Shanghai, China. [6]Shanghai Matwings Technology Co., Ltd., Shanghai, China. [7]Shanghai Artificial Intelligence Laboratory, Shanghai, China.

References

1. Preuer, K., Renz, P., Unterthiner, T., Hochreiter, S., Klambauer, G.: Fréchet chemnet distance: A metric for generative models for molecules in drug discovery. Journal of chemical information and modeling **58 9**, 1736–1741 (2018)
2. Olivecrona, M., Blaschke, T., Engkvist, O., Chen, H.: Molecular de-novo design through deep reinforcement learning. Journal of Cheminformatics **9**(1) (2017). doi:10.1186/s13321-017-0235-x
3. Segler, M.H.S., Kogej, T., Tyrchan, C., Waller, M.P.: Generating focused molecule libraries for drug discovery with recurrent neural networks. ACS Central Science **4**(1), 120–131 (2018). doi:10.1021/acscentsci.7b00512
4. Li, Y., Zhang, L., Liu, Z.: Multi-objective de novo drug design with conditional graph generative model. Journal of Cheminformatics **10**(1) (2018). doi:10.1186/s13321-018-0287-6
5. Zhavoronkov, A., Ivanenkov, Y.A., Aliper, A., Veselov, M.S., Aladinskiy, V.A., Aladinskaya, A.V., Terentiev, V.A., Polykovskiy, D.A., Kuznetsov, M.D., Asadulaev, A., Volkov, Y., Zholus, A., Shayakhmetov, R.R., Zhebrak, A., Minaeva, L.I., Zagribelnyy, B.A., Lee, L.H., Soll, R., Madge, D., Xing, L., Guo, T., Aspuru-Guzik, A.: Deep learning enables rapid identification of potent ddr1 kinase inhibitors. Nature Biotechnology **37**(9), 1038–1040 (2019). doi:10.1038/s41587-019-0224-x
6. Shi, C., Xu, M., Zhu, Z., Zhang, W., Zhang, M., Tang, J.: Graphaf: a flow-based autoregressive model for molecular graph generation. ArXiv **abs/2001.09382** (2020)
7. Blaschke, T., Arús-Pous, J., Chen, H., Margreitter, C., Tyrchan, C., Engkvist, O., Papadopoulos, K., Patronov, A.: Reinvent 2.0: An ai tool for de novo drug design. Journal of Chemical Information and Modeling **60**(12), 5918–5922 (2020). doi:10.1021/acs.jcim.0c00915
8. Yang, Y., Zheng, S., Su, S., Zhao, C., Xu, J., Chen, H.: Syntalinker: automatic fragment linking with deep conditional transformer neural networks. Chemical Science **11**(31), 8312–8322 (2020). doi:10.1039/D0SC03126G
9. Wang, J., Hsieh, C.-Y., Wang, M., Wang, X., Wu, Z., Jiang, D., Liao, B., Zhang, X., Yang, B., He, Q., Cao, D., Chen, X., Hou, T.: Multi-constraint molecular generation based on conditional transformer, knowledge distillation and reinforcement learning. Nature Machine Intelligence **3**(10), 914–922 (2021). doi:10.1038/s42256-021-00403-1
10. Kadurin, A., Aliper, A., Kazennov, A., Mamoshina, P., Vanhaelen, Q., Khrabrov, K., Zhavoronkov, A.: The cornucopia of meaningful leads: Applying deep adversarial autoencoders for new molecule development in oncology. Oncotarget **8**(7) (2016)
11. Gómez-Bombarelli, R., Wei, J.N., Duvenaud, D., Hernández-Lobato, J.M., Sánchez-Lengeling, B., Sheberla, D., Aguilera-Iparraguirre, J., Hirzel, T.D., Adams, R.P., Aspuru-Guzik, A.: Automatic chemical design using a data-driven continuous representation of molecules. ACS Central Science **4**(2), 268–276 (2018). doi:10.1021/acscentsci.7b00572
12. Prykhodko, O., Johansson, S.V., Kotsias, P.-C., Arús-Pous, J., Bjerrum, E.J., Engkvist, O., Chen, H.: A de novo molecular generation method using latent vector based generative adversarial network. Journal of Cheminformatics **11**(1) (2019). doi:10.1186/s13321-019-0397-9
13. Zheng, S., Yan, X., Gu, Q., Yang, Y., Du, Y., Lu, Y., Xu, J.: Qbmg: quasi-biogenic molecule generator with deep recurrent neural network. Journal of Cheminformatics **11**(1) (2019). doi:10.1186/s13321-019-0328-9
14. Blaschke, T., Engkvist, O., Bajorath, J., Chen, H.: Memory-assisted reinforcement learning for diverse molecular de novo design. Journal of Cheminformatics **12**(1) (2020). doi:10.1186/s13321-020-00473-0
15. Grisoni, F., Moret, M., Lingwood, R., Schneider, G.: Bidirectional molecule generation with recurrent neural networks. Journal of Chemical Information and Modeling **60**(3), 1175–1183 (2020). doi:10.1021/acs.jcim.9b00943
16. Kuznetsov, M., Polykovskiy, D.: Molgrow: A graph normalizing flow for hierarchical molecular generation. ArXiv **abs/2106.05856** (2021)
17. Spiegel, J.O., Durrant, J.D.: Autogrow4: an open-source genetic algorithm for de novo drug design and lead optimization. Journal of Cheminformatics **12**(1) (2020). doi:10.1186/s13321-020-00429-4
18. Boitreaud, J., Mallet, V., Oliver, C., Waldispühl, J.: Optimol: Optimization of binding affinities in chemical space for drug discovery. Journal of Chemical Information and Modeling **60**(12), 5658–5666 (2020). doi:10.1021/acs.jcim.0c00833
19. Zhou, Z., Kearnes, S., Li, L., Zare, R.N., Riley, P.: Optimization of molecules via deep reinforcement learning. Scientific Reports **9**(1) (2019). doi:10.1038/s41598-019-47148-x

20. Li, Y., Pei, J., Lai, L.: Structure-based de novo drug design using 3d deep generative models. Chem. Sci. **12**, 13664–13675 (2021). doi:10.1039/D1SC04444C
21. Peng, X., Luo, S., Guan, J., Xie, Q., Peng, J., Ma, J.: Pocket2mol: Efficient molecular sampling based on 3d protein pockets. In: International Conference on Machine Learning (2022)
22. Ma, B., Terayama, K., Matsumoto, S., Isaka, Y., Sasakura, Y., Iwata, H., Araki, M., Okuno, Y.: Structure-based de novo molecular generator combined with artificial intelligence and docking simulations. Journal of Chemical Information and Modeling **61**(7), 3304–3313 (2021). doi:10.1021/acs.jcim.1c00679
23. Baig, M.H., Ahmad, K., Roy, S., Ashraf, J.M., Adil, M., Siddiqui, M.H., Khan, S., Kamal, M.A., Provazník, I., Choi, I.: Computer aided drug design: Success and limitations. Current pharmaceutical design **22 5**, 572–81 (2016)
24. Gaulton, A., Bellis, L.J., Bento, A.P., Chambers, J., Davies, M., Hersey, A., Light, Y., McGlinchey, S., Michalovich, D., Al-Lazikani, B., Overington, J.P.: Chembl: a large-scale bioactivity database for drug discovery. Nucleic Acids Research **40**(D1), 1100–1107 (2012). doi:10.1093/nar/gkr777
25. Zhang, C.-H., Stone, E.A., Deshmukh, M., Ippolito, J.A., Ghahremanpour, M.M., Tirado-Rives, J., Spasov, K.A., Zhang, S., Takeo, Y., Kudalkar, S.N., Liang, Z., Isaacs, F.J., Lindenbach, B., Miller, S.J., Anderson, K.S., Jorgensen, W.L.: Potent noncovalent inhibitors of the main protease of sars-cov-2 from molecular sculpting of the drug perampanel guided by free energy perturbation calculations. ACS Central Science **7**, 467–475 (2021)
26. Wang, Z., Sun, H., Yao, X., Li, D., Xu, L., Li, Y., Tian, S., Hou, T.: Comprehensive evaluation of ten docking programs on a diverse set of protein–ligand complexes: the prediction accuracy of sampling power and scoring power. Phys. Chem. Chem. Phys. **18**, 12964–12975 (2016). doi:10.1039/C6CP01555G
27. Chung, J., Gülçehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. ArXiv **abs/1412.3555** (2014)
28. Sunseri, J., Koes, D.R.: Pharmit: interactive exploration of chemical space. Nucleic Acids Research **44**(W1), 442–448 (2016). doi:10.1093/nar/gkw287
29. Schrödinger, LLC: The PyMOL Molecular Graphics System, Version 1.8 (2015)
30. Polykovskiy, D., Zhebrak, A., Sánchez-Lengeling, B., Golovanov, S., Tatanov, O., Belyaev, S., Kurbanov, R., Artamonov, A.A., Aladinskiy, V., Veselov, M., Kadurin, A., Nikolenko, S.I., Aspuru-Guzik, A., Zhavoronkov, A.: Molecular sets (moses): A benchmarking platform for molecular generation models. Frontiers in Pharmacology **11** (2018)
31. Morgan, H.L.: The generation of a unique machine description for chemical structures-a technique developed at chemical abstracts service. Journal of Chemical Documentation **5**, 107–113 (1965)
32. Vankadara, S., Wong, Y.X., Liu, B., See, Y.Y., Tan, L.H., Tan, Q.W., Wang, G., Karuna, R., Guo, X., Tan, S.T., Fong, J.Y., Joy, J., Chia, C.S.B.: A head-to-head comparison of the inhibitory activities of 15 peptidomimetic sars-cov-2 3clpro inhibitors. Bioorganic & Medicinal Chemistry Letters **48**, 128263 (2021). doi:10.1016/j.bmcl.2021.128263

Figures

[width=16cm]figures/fig1.pdf

Figure 1 The concept of molecular generative models, and the architecture of LS-MolGen approach. (a) The concept of ligand-based method, (b) structure-based method, and (c) ligand-and-structure-based method. Illustration in a, b, and c, the green region represents the available chemical space, the yellow region represents the chemical space of known bioactive ligands, and the blue arrow represents the exploration of chemical space. (d) The pipeline of the LS-MolGen approach. (e) The architecture of the prior model with RNN. (f) The illustration of transfer learning model. (g) Exploration of chemical space of reinforcement learning combined with molecular docking.

[width=16cm]figures/fig2.pdf

Figure 2 Quantitative evaluations of LS-MolGen. (a) Evolution of the explored molecular average score in the reinforcement learning loop. (b) Chemical space of generated molecules and bioactive ligands of EGFR visualized by t-SNE dimensionality reduction. (c) Distribution of molecular docking score to EGFR. (d-f) Distributions of properties, including molecular weight, QED, and SA score.

Tables

Table 1 The evaluation metrics of molecules generated by the baseline models and LS-MolGen. Each model generated 5,000 molecules for six times independently, and 30,000 molecules were generated from each model eventually.

Model		Validity	Uniqueness	Diversity	Novelty	Recovery	Active rate	Success rate
L-based[a]	AAE [10]	0.950 ± 0.003	0.523 ± 0.005	0.248 ± 0.004	0.104 ± 0.004	0.621 ± 0.014	0.869 ± 0.004	0.131 ± 0.004
	VAE [11]	0.988 ± 0.002	0.426 ± 0.005	0.180 ± 0.004	0.043 ± 0.002	0.828 ± 0.012	0.897 ± 0.004	0.082 ± 0.003
	LatentGAN [12]	0.731 ± 0.004	0.697 ± 0.003	0.655 ± 0.011	0.591 ± 0.011	0.134 ± 0.002	0.642 ± 0.004	0.505 ± 0.005
S-based[b]	REINVENT [2]	0.992 ± 0.021	0.990 ± 0.020	0.720 ± 0.158	0.595 ± 0.159	0.134 ± 0.030	0.904 ± 0.234	0.826 ± 0.237
	Pocket2Mol [21]	1.000 ± 0.000	1.000 ± 0.000	0.488 ± 0.043	0.475 ± 0.043	0.092 ± 0.041	0.778 ± 0.175	0.778 ± 0.176
	MolDQN [19]	1.000 ± 0.000	1.000 ± 0.000	0.415 ± 0.007	0.414 ± 0.007	0.005 ± 0.000	0.495 ± 0.005	0.495 ± 0.005
L-S-based[c]	SBMolGen [22]	1.000 ± 0.000	0.999 ± 0.000	0.867 ± 0.002	0.687 ± 0.003	0.070 ± 0.001	0.351 ± 0.002	0.346 ± 0.002
	LS-MolGen	1.000 ± 0.000	0.997 ± 0.005	$\mathbf{0.979 \pm 0.007}$	$\mathbf{0.966 \pm 0.015}$	0.202 ± 0.037	$\mathbf{0.999 \pm 0.002}$	$\mathbf{0.950 \pm 0.009}$

[a] L-based denotes ligand-based models; [b] S-based denotes structure-based models; [c] L-S-based denotes ligand-and-structure-based models.

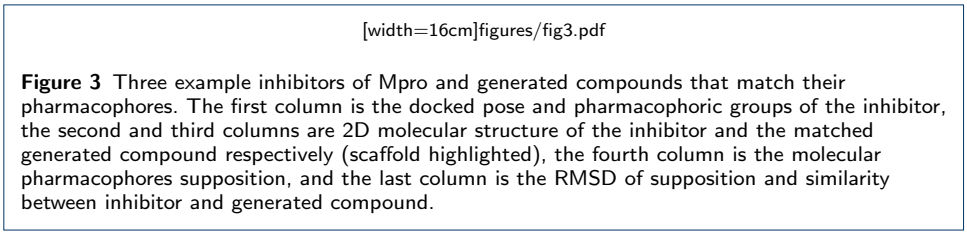

Figure 3 Three example inhibitors of Mpro and generated compounds that match their pharmacophores. The first column is the docked pose and pharmacophoric groups of the inhibitor, the second and third columns are 2D molecular structure of the inhibitor and the matched generated compound respectively (scaffold highlighted), the fourth column is the molecular pharmacophores supposition, and the last column is the RMSD of supposition and similarity between inhibitor and generated compound.

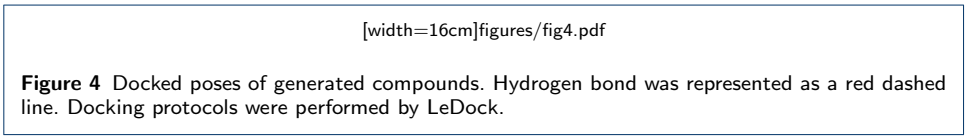

Figure 4 Docked poses of generated compounds. Hydrogen bond was represented as a red dashed line. Docking protocols were performed by LeDock.

Table 2 Performance Comparison of Different Model Variants in EGFR Ablation Experiment

Model	Validity	Uniqueness	Diversity	Novelty	Recovery	Active rate	Success rate
Prior	0.936	0.936	0.882	0.564	0.088	0.421	0.338
Prior+Transfer	0.895	0.890	0.376	0.191	**0.506**	0.811	0.150
Prior+Agent	**1.000**	**0.998**	**0.981**	**0.974**	0.158	0.901	0.857
LS-MolGen	**1.000**	0.997	0.979	0.966	0.203	**1.000**	**0.987**

Figures

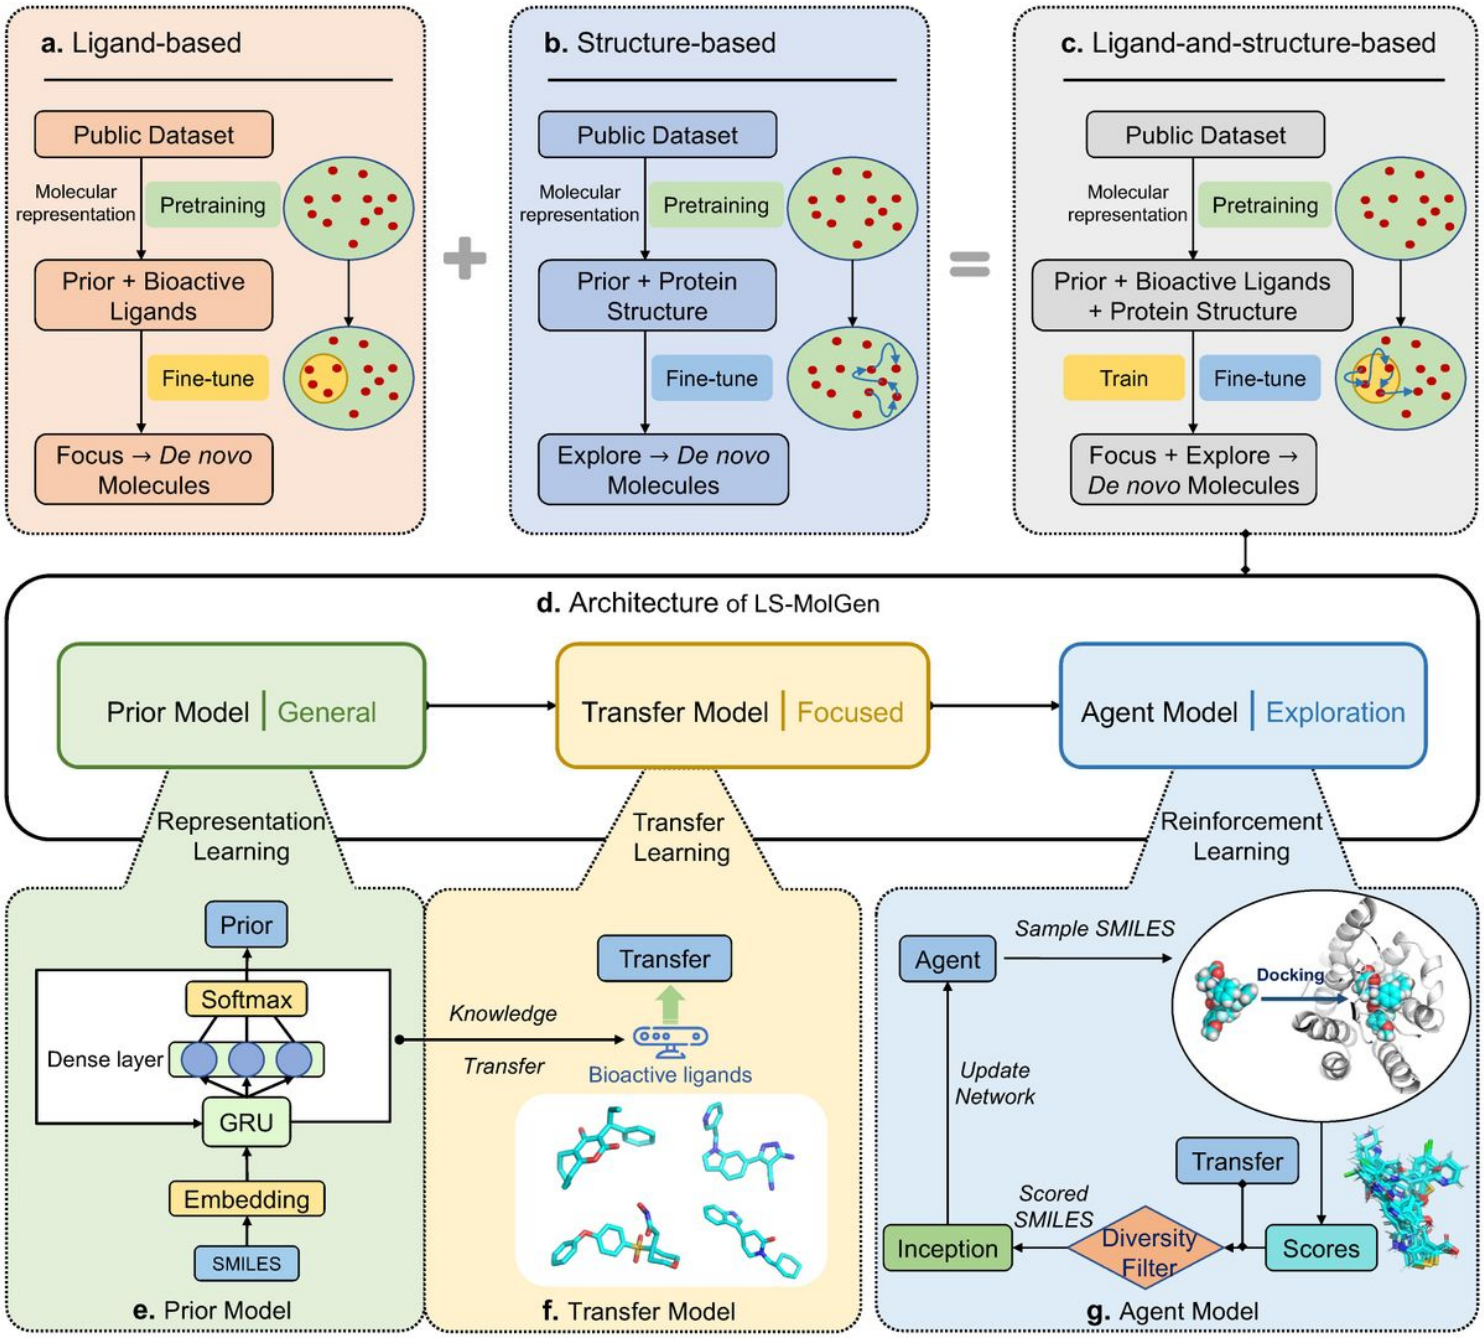

Figure 1

The concept of molecular generative models, and the architecture of LS-MolGen approach. (a) The concept of ligand-based method, (b) structure-based method, and (c) ligand-and-structure-based method. Illustration in a, b, and c, the green region represents the available chemical space, the yellow region represents the chemical space of known bioactive ligands, and the blue arrow represents the exploration of chemical space. (d) The pipeline of the LS-MolGen approach. (e) The architecture of the prior model with RNN. (f) The illustration of transfer learning model. (g) Exploration of chemical space of reinforcement learning combined with molecular docking.

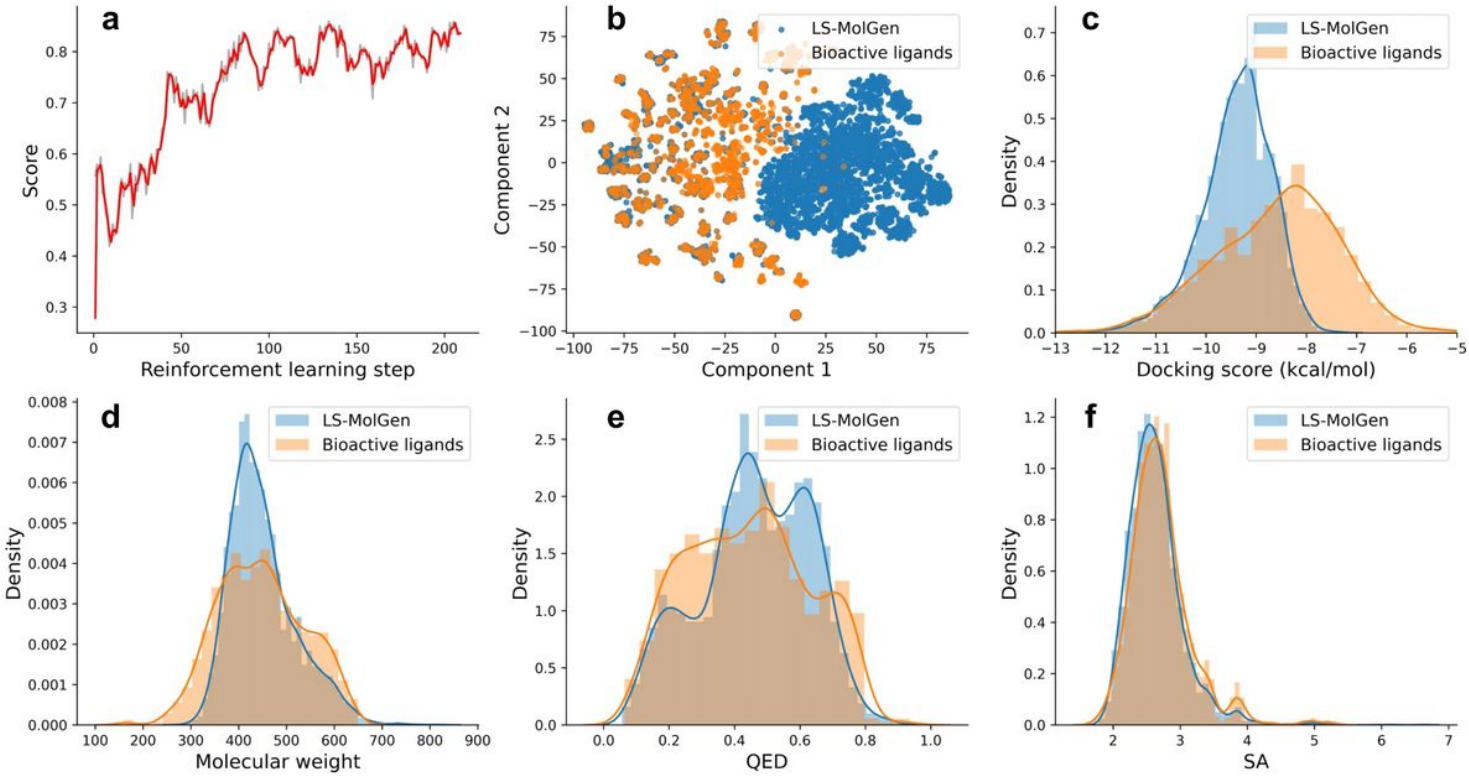

Figure 2

Quantitative evaluations of LS-MolGen. (a) Evolution of the explored molecular average score in the reinforcement learning loop. (b) Chemical space of generated molecules and bioactive ligands of EGFR visualized by t-SNE dimensionality reduction. (c) Distribution of molecular docking score to EGFR. (d-f) Distributions of properties, including molecular weight, QED, and SA score.

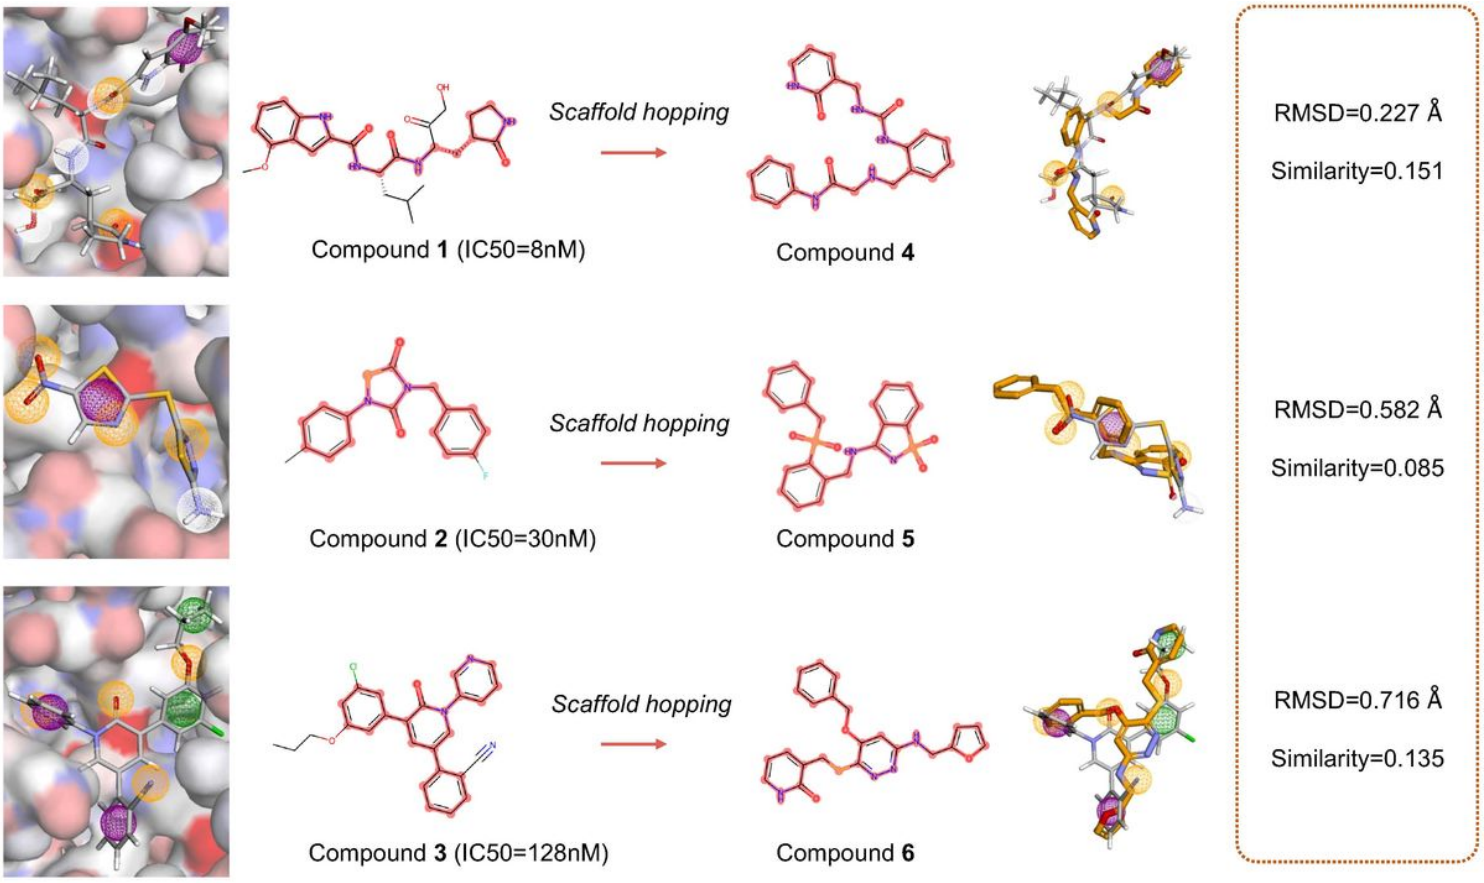

Figure 3

Three example inhibitors of Mpro and generated compounds that match their pharmacophores. The first column is the docked pose and pharmacophoric groups of the inhibitor, the second and third columns are 2D molecular structure of the inhibitor and the matched generated compound respectively (scaffold highlighted), the fourth column is the molecular pharmacophores supposition, and the last column is the RMSD of supposition and similarity between inhibitor and generated compound.

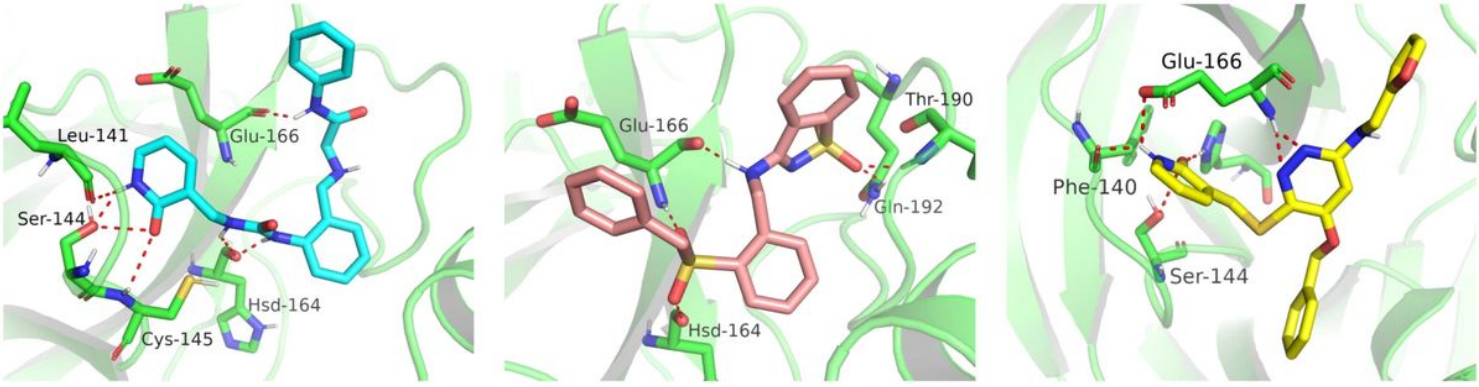

Docking score = -9.71 kcal/mol

Compound **4**

Docking score = -9.02 kcal/mol

Compound **5**

Docking score = -9.45 kcal/mol

Compound **6**

Figure 4

Docked poses of generated compounds. Hydrogen bond was represented as a red dashed line. Docking protocols were performed by LeDock.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supportinginformation.pdf